Chapter 7

Fabry-Perot Cavities for Fun and Profit

7.1 REO/ATF Mirrors at the Current State of the Art

If the cavity is the heart of a cavity QED experiment, the highly reflective mirrors are the heart of the cavity. We obtain these mirrors from the group of Ramin Lalezari at Research Electro-Optics (REO) in Boulder, CO – or within the last year from the same individual in his new effort at Advanced Thin Films (ATF), also near Boulder. The substrates, currently BK7 glass or fused silica, are superpolished on front (HR) and back surfaces. A standard antireflection coating is applied to the back surface. The mirror surface is coated with a multilayer stack of dielectric materials so that coating layers have alternating high and low index of refraction; the layer thickness is $\lambda_c/4$ for a coating centered at λ_c.

These coating techniques are capable of producing mirror transmission T at or below 10^{-6} (1 ppm). However, current technology has yet to push mirror loss A below one to a few ppm. Thus cavity finesses in the range of $F \approx 10^5$ to 10^6 constitute the current state of the art for high-reflectivity, low-loss mirrors and coatings. These mirrors are useful not only in cavity QED but in numerous other scientific applications that exploit the high sensitivity they offer [105]. Mirror absorption/scatter losses set a limit on F and are furthermore a hindrance to signal extraction when nonlinear

interactions are present. They are a critical limiting factor in the loss rate for present cavity QED systems; for the very short cavities used in these experiments, loss rates associated with A are usually similar in size to the atomic spontaneous emission rates. To build robust quantum computing/communications devices from cavity QED components, it will be necessary to improve the ratio of mirror transmission (useful information) to mirror losses (loss of coherence).

In addition to the ratio of transmission to loss, another significant consideration is the role of mirror dispersion in determining cavity properties. At the current very short cavity lengths (about 20 half-wavelengths of the optical field), it is important to consider that the standing-wave light field inside a cavity penetrates into the mirror coatings, giving a larger mode volume than would be otherwise expected from the physical distance between the mirror surfaces. When the micro-cavities are pushed to shorter lengths in the quest for greater interaction strengths, the leakage field into the mirror coatings will have a non-negligible effect on the cavity mode structure.

A detailed coating model of the current experimentally-employed mirrors is described in [18], and a manual on the cleaning and evaluation of these mirrors can be found in [19]. Mount design considerations were presented in Chapter 6. In this chapter I focus on optical, geometric, and cavity QED properties of the cavities themselves.

7.2 Birefringence Issues

One unavoidable consequence of using high-finesse optical cavities seems to be a nagging issue of cavity birefringence. In a nutshell, the problem is that with cavity finesse in the range of $10^5 < F < 10^6$ or higher, any birefringent phase shift δ per round trip is ultimately enhanced by the very large finesse.

A very quick computation calculates the splitting between cavity resonance frequencies (ν_a, ν_b) for TEM_{00} modes of orthogonal linear polarizations. Assume $\hat{a}$ acquires a phase shift δ relative to $\hat{b}$ on each round trip in the cavity. Then the resonance conditions for a given longitudinal mode imply $\frac{2\pi}{\lambda_a}2l + \delta = \frac{2\pi}{\lambda_b}2l$ and so $\frac{2\pi l}{c}(\nu_b - \nu_a) =$

$\frac{2\pi l}{c}\Delta\nu = \delta/2$. Thus $\Delta\nu = \frac{\delta}{2\pi}\frac{c}{2l}$ and substituting $\frac{c}{2l} = FSR = (FWHM)(F)$, we have the result

$$\frac{\Delta\nu}{FWHM} = \frac{\delta}{2\pi}F \tag{7.1}$$

for the birefringent splitting relative to the cavity width. Thus even phase shifts $\delta \sim 10^{-7}$ are quite noticeable; $\delta \sim 10^{-6}$, for instance, will completely scramble circular polarization injected into an otherwise beautiful high-finesse cavity.

Notice an equivalent formulation of this result is that, if the cavity is resonant with one polarization, the other polarization is detuned $\Delta\nu$ off its resonance. Thus any orthogonally polarized light that is transmitted gets a phase shift of $tan^{-1}(2 * detuning/FWHM)$ [106], i.e., $\theta = tan^{-1}(\frac{2\Delta\nu}{FWHM}) = tan^{-1}(\delta\frac{F}{\pi})$. For small overall phase shifts θ, this becomes $\theta \approx \delta\frac{F}{\pi}$, meaning the *single-pass* phase shift $\delta/2$ is simply magnified by the standard cavity enhancement factor $2F/\pi$.

If the previous two paragraphs were a long way of saying that even a small phase shift δ is evil, where does the evil arise? From a variety of circumstantial evidence over the years, the most likely culprit is stress-induced birefringence in the dieletric coating stack of the mirrors. This conjecture is discussed in Chapter 6 and in Christina Hood's thesis [19], but a short summary is that cavities tend to be more birefringent when the mirrors are glued or clamped down more firmly and/or at positions closer to the coated surface. Often they become birefringent when they are glued and then baked, a process that could induce stress on the coatings as glue hardens and pulls the mirrors. Additionally, I have some experience with 7.75 mm substrates which are clamped with a set screw close to the coated surface; cavities made from two of these tend to be least birefringent when the mirrors are well centered on each other and thus the mode stays away from the point of greatest stress.

Another, but weaker, candidate for our suspicions is some sort of polarization-dependent absorption, scattering, or reflectivity resulting from defects or irregularities of wavelength scale or smaller. These effects would lead to different cavity decay widths for orthogonal linear polarizations, but would only cause a peak splitting $\Delta\nu$ if they altered the spatial mode profile and thus induced different Guoy phase shifts

for the modes (see, e.g., [106]). This possibility should be easy to distinguish from the stress hypothesis by comparing the peak splitting to the difference in peak width for the two cavity eigenpolarizations. The evidence indicates stress is to blame in all but one instance of cavity birefringence measured over six and a half years. Thus I will concentrate on the situation where $\hat{a}$ and $\hat{b}$ are birefringent axes with equal cavity linewidths.

There are many possible measures of birefringence, and any one will suffice for the purpose of comparing different cavities to one another. Care is essential, though, when relating different measurement methods, since the easiest few measurements fit together in simple yet not-so-simple ways. Much of the difficulty lies in the typical measurement process of scanning over the cavity resonance(s) and doing a polarization-sensitive measurement of transmitted peak height. This process is well-suited to a tabletop building/testing situation where the cavity length is not well stabilized. However, because the lineshape too is polarization sensitive, this measurement is not as clean and direct as one could hope. Here I have looked at the three or four measurements most common in our group over last few years, relating each of them to one another and to the underlying birefringent phase shift δ per round trip.

The general technique for most measurements is depicted in Figure 7.1. The cavity length is scanned over the entire resonance profile; light is injected into the cavity after some set of polarization optics, transmitted through the cavity, and analyzed on some (different) set of output polarization optics plus a photodetector. For interpretation and the discussion that follows, we assume the time to scan over any resonance feature is itself much greater than the cavity buildup/ringdown time, so that the transmitted lineshape unambiguously reflects steady-state cavity transmission as a function of length.

- Perhaps the simplest way to measure a birefringence is to directly measure a peak splitting as a fraction of underlying peak width. This is simple if the peaks are well-resolved, $\Delta\nu \gtrsim FWHM$. A cavity linewidth is measured for on-axis linear input polarization, and a peak splitting is read directly off the same scan

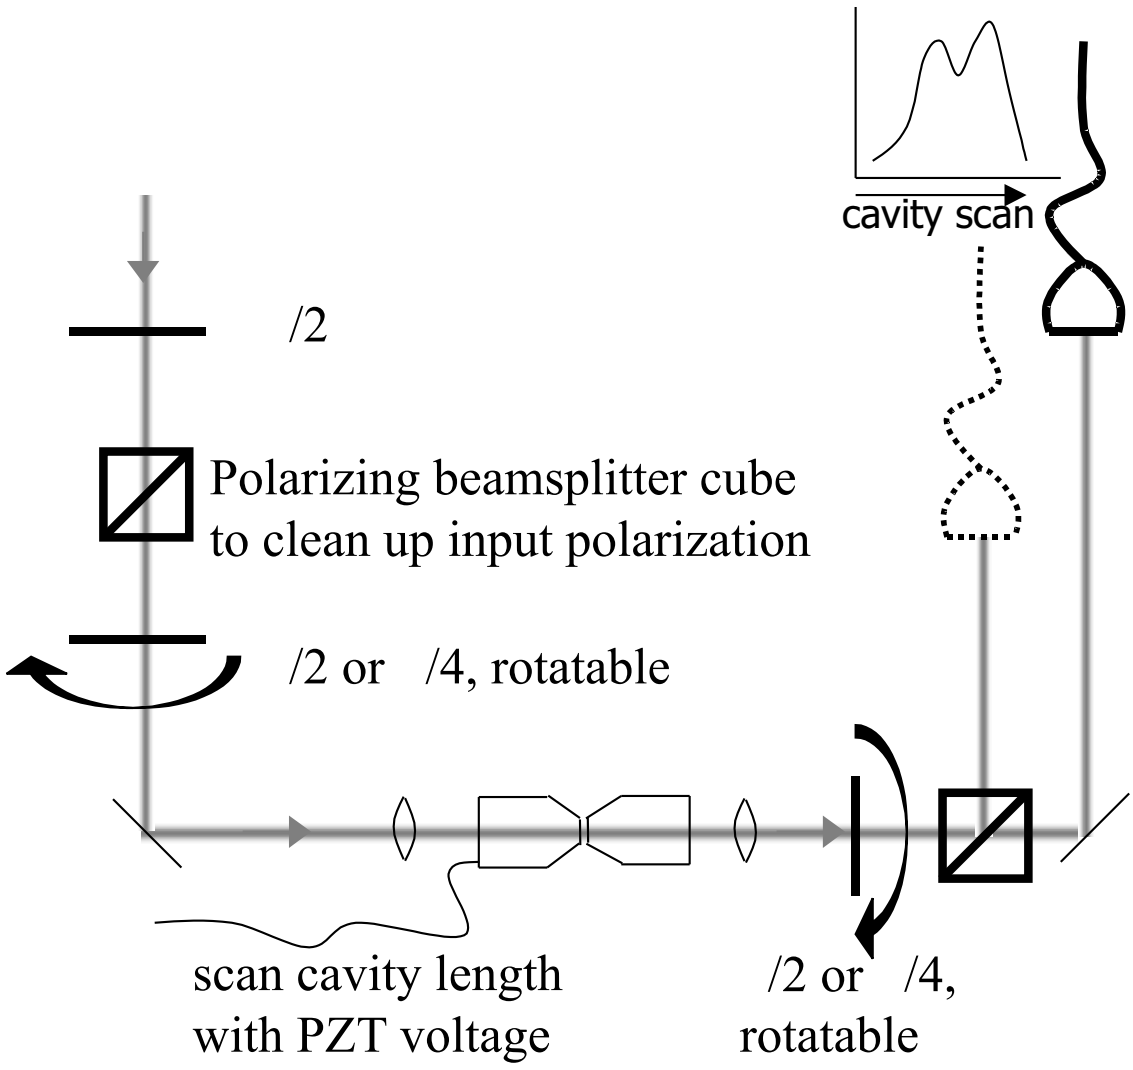

Figure 7.1: Setup for measurements of cavity birefringence

for circular input polarization (or off-axis linear input). If $\Delta\nu \gg FWHM$, the peak splitting is just $\Delta\nu$. If the splitting is smaller, some care must be taken in backing out $\Delta\nu$ from the summed intensity transmission of the two modes.

For somewhat smaller birefringence (line not resolved into two peaks), one can measure the width separately on a birefringent axis and halfway between axes; peak splitting can always in theory be inferred from this. In practice, for splittings smaller than $\Delta\nu \sim FWHM$, direct peak splitting measurements may be hard to perform with any reasonable accuracy even though the birefringence is still large enough to be of interest.

- For small birefringence, a fairly simple measurement for diagnostic purposes is to scan the cavity, inject linearly polarized light, and detect it on a rotatable linear polarizer at the output (e.g., with a half waveplate and polarizing beamsplitter cube in front of the photodetector). On-axis or for no birefringence, the output light remains linear and one can rotate the output selector to transmit all of it or cancel it entirely (Figure 7.2(a)). Off-axis, though, the transmitted field of one polarization is phase shifted relative to the other; the overall output polarization is not linear but varies across the (split) cavity line. At overall line center, halfway between ν_a and ν_b, the polarizations have relative phase shift $2tan^{-1}(\Delta\nu/FWHM)$. We inject light off-axis along $(\hat{a}+\hat{b})/\sqrt{2}$. For small enough birefringence, rotating an output linear polarizer gives maximum transmitted peak height at $(\hat{a}+\hat{b})/\sqrt{2}$, and minimum peak height at the orthogonal setting. (See Figure 7.2(b).) For larger birefringences, there will be maxima slightly away from the original axis due to peak splitting effects. Thus as long as the transmitted peak height achieves maximum and minimum values at orthogonal output analyzer settings ($\theta_{out} = (\pi/4, 3\pi/4)$, we can record the max and min heights and apply

$$Contrast_{lin} = \frac{max - min}{max + min} = cos(2tan^{-1}(\Delta\nu/FWHM)) = \frac{1 - (\Delta\nu/FWHM)^2}{1 + (\Delta\nu/FWHM)^2} \tag{7.2}$$

As soon as the strict $\pi/2$ max/min condition is violated (Figure 7.2(c,d,e)), the formula above no longer applies. $\Delta\nu$ can still be extracted, but it is best to go to a different measurement method anyway. Thus this measurement works nicely for $\Delta\nu \lesssim 0.42(FWHM)$ or $Contrast_{lin} \gtrsim 0.70$.

- A similar measurement involves injecting circular polarization and measuring transmitted peak heights through an output linear polarizer. In this case there should be no dependence on analyzer setting in the perfect case, and contrast grows with birefringence. As long as contrast goes from max to min in $\pi/2$ rotation of the output polarizer (Figure 7.3(a,b,c)), max and min occur at line center and we have

$$Contrast_{circ} = \frac{max - min}{max + min} = sin(2tan^{-1}(\Delta\nu/FWHM)) = \frac{2\Delta\nu/FWHM}{1 + (\Delta\nu/FWHM)^2}. \tag{7.3}$$

  Just as with the previous measurement, this no longer holds when the $\pi/2$ condition is violated (Figure 7.3(d,e)) and it is best to seek another measurement. However, it turns out this measurement holds a bit further, out to $\Delta\nu \lesssim 0.64(FWHM)$ or $Contrast_{circ} \lesssim 0.44$.

- Other measurements are thus required for intermediate birefringence where the peaks are not yet clearly resolved but the two methods above already give dubious results. A nice method is to inject circular or off-axis linear polarization and measure the output transmission profile through a selector for the original polarization. When the birefringence is still too small to resolve the overall peaks, this polarization-sensitive measurement will show two resolved peaks from which the underlying splitting can be inferred (Figure 7.4).

- Yet another intermediate-splitting method, employed by Jason McKeever and Dan Stamper-Kurn, is to inject circular polarization, rotate a linear analyzer, and watch the transmission peak move around line center. They recorded a "maximum" peak height and the smaller peak height when the line was centered. A measurement of the splitting can be backed out of these quantities

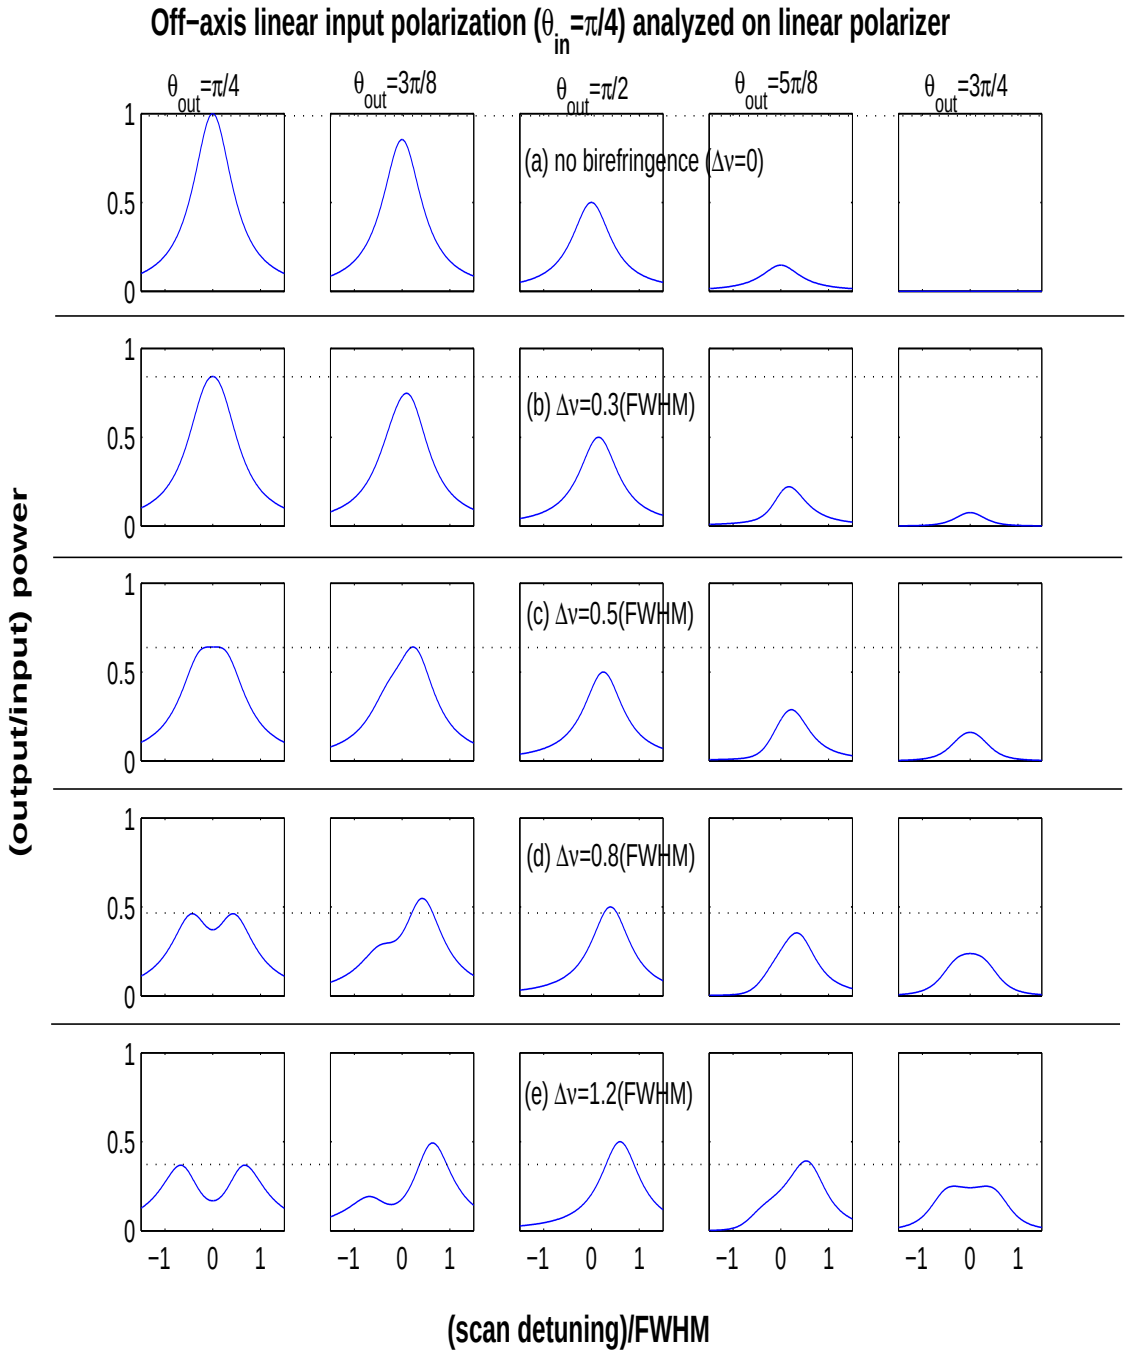

Figure 7.2: Calculated transmission of off-axis linear input light ($\theta_{in} = \pi/4$, or $\frac{1}{\sqrt{2}}(\hat{a}+\hat{b})$) through cavity and output linear polarizer for five different values of cavity birefringence

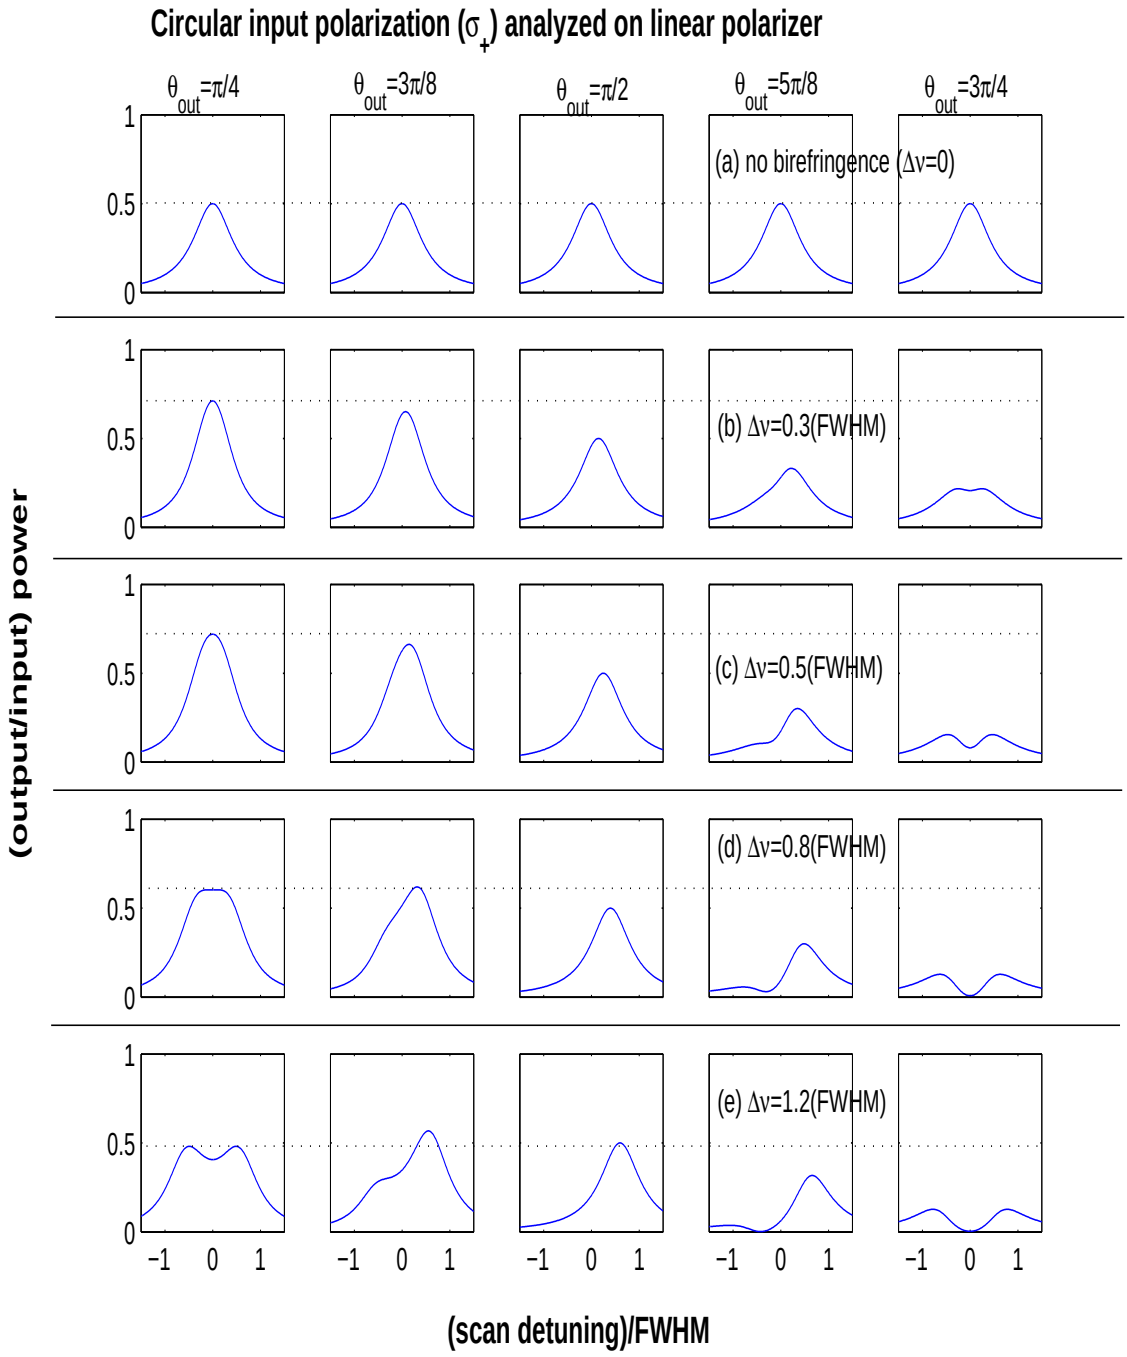

Figure 7.3: Calculated transmission of circular input light (σ_+, or $\frac{1}{\sqrt{2}}(\hat{a}+i\hat{b})$) through cavity and output linear polarizer for five different values of cavity birefringence.

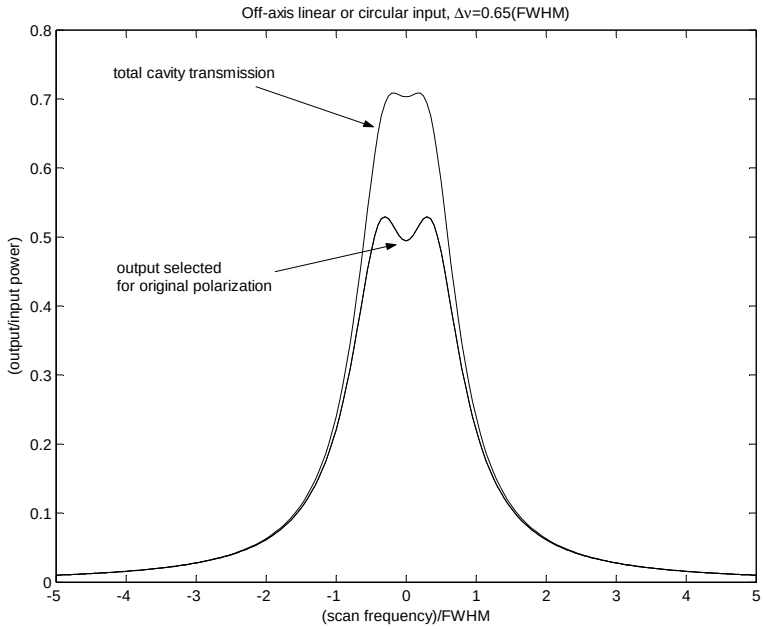

Figure 7.4: Preservation of input linear or circular polarization through birefringent cavity.

in the intermediate regime. However, the contrast on this measurement can be improved significantly by doing the same thing but with off-axis linear input light. Aside from the trouble needed to identify the axes, the linear input measurement is definitely superior. (Refer to Figures 7.2 and 7.3.)

- Often the "real" question from the operative point of view is: if we drive the cavity with a given polarization containing both $\hat{a}$ and $\hat{b}$, how much does the actual intracavity polarization differ from our drive? This question is crucial for quantitatively predicting scattering rates and all atom-field interactions in the cavity. Unfortunately the question is not a simple one. It depends on the cavity detuning from resonance with each probe polarization. If the probe in question is the strongly-coupled cavity QED light, the polarization is also altered by the presence of an atom, since the atom itself shifts system resonances. Thus such a measurement can be performed for the empty cavity, but in reality any carefully set input polarization will shift as an atom enters the cavity and couples strongly to the light.

7.3 What Cavities Are Good for What?

I have attempted to collect in one place the information that must be considered when designing a new physics cavity (i.e., selecting mirror reflectivity, radius of curvature, cavity length, etc.). Much of this material is common knowledge, comes from simple geometry, or is written up elsewhere. Nevertheless, my hope is that by including the old with the new I can provide a single primary reference on this subject.

Several quite different physics considerations drive cavity design. In the experiments we have done for the past several years, atoms are dropped into the cavity from a MOT, so geometry determines the flux of atoms through the cavity mode and also their phase space distribution (particularly fall velocity and velocity along the cavity axis). Most likely, light besides the near-resonant probe is desired inside the cavity volume, so it is important to consider what longitudinal and transverse modes the cavity will support – as well as ease of coupling light in, detecting what comes out, or even focusing it through from the side. Finally, we come to the explicitly cavity-QED considerations; values of g_0 and κ are certainly crucial, but for each experimental design the relevant figure of merit may be slightly different. I discuss important ratios and sensitivities for a couple of purposes to give a flavor.

I will devote separate subsections to "atom" factors, "light" factors, and "cavity QED" factors. First, however, I catalog basic relations of cavity geometry and mode structure that determine the quantities to be discussed throughout the later sections.

7.3.1 Basic Cavity Geometry and Mode Structure

We consider a Fabry-Perot cavity of length l made from spherical mirrors shown in Figure 7.5. Each mirror has radius of curvature R, diameter d, and transmission and losses (T,A). If the two mirrors are different the subscripts $(1, 2)$ are appended to each of these quantities. We treat lengths $l \ll R$.

The physical length of the cavity is $l_{phys} = n\lambda/2$. This is the separation between the two mirror surfaces at mode center. Measurements of cavity mode spacing and related quantities yield a somewhat different l_{eff} because of nonzero intensities in

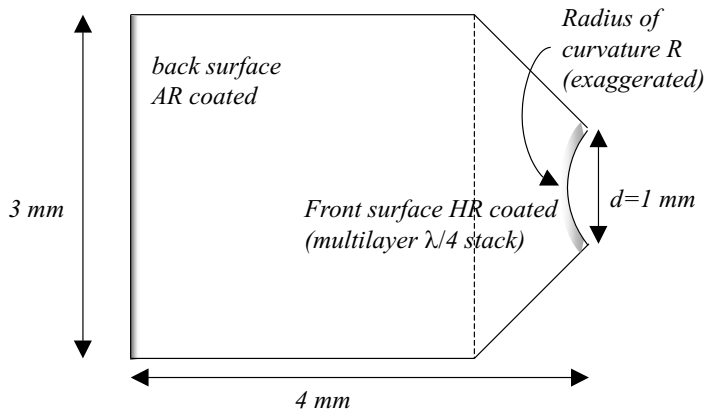

Figure 7.5: High reflectivity, low-loss mirrors for cavity QED; side view of cylinder-plus-cone substrate shape.

the first few layers of mirror coating. For mirrors in the T95 run, for example, $l_{eff} = (n + 1.63)\lambda/2$. More on this subject can be found in the next section and in [18].

Since the mirrors are spherical, there is an indentation from mirror edge to center, and thus the *gap* between mirror edges is less than l_{phys}. The indentation of each mirror is $d^2/8R$, giving

$$gap = l_{phys} - (d_1^2/8R_1 + d_2^2/8R_2). \tag{7.4}$$

The Gaussian waist of the cavity mode at wavelength λ is given by

$$w_0 = (\frac{l_{phys}}{2}(R - l_{phys}/2)(\lambda/\pi)^2)^{1/4} \approx (\frac{1}{2}(\frac{\lambda}{\pi})^2 R l_{phys})^{1/4}. \tag{7.5}$$

for a symmetric cavity. Note that the cavity mode volume therefore scales as $w_0^2 l_{eff} \propto R^{1/2} l_{phys}^{1/2} l_{eff}$. Generally $l_{phys} \approx l_{eff} \approx l$, so we may write this scaling as $R^{1/2} l^{3/2}$.

7.3.2 Atoms in the Cavity

Atoms are collected in a MOT as close as possible to the cavity, sub-Doppler cooled, and dropped onto the mirror gap. For fixed MOT properties and sub-Doppler temperature, cavity geometry will determine the flux of atoms through the cavity mode

volume and their position/velocity distribution. How does the flux of atoms through the mode depend on cavity geometry? We imagine the MOT dropping unobstructed to the top edge of the mirror surfaces, and begin our considerations then:

- 1) Initial x (axial position) must be within the gap, so atom flux Φ scales as gap.

- 2) For each initial x, geometry allows

$$\frac{-l_{phys}/2 - x}{d/2} \leq v_x/v_z \leq \frac{l_{phys}/2 - x}{d/2} \tag{7.6}$$

  in order for an atom to arrive at the cavity mode before hitting a mirror. The vertical velocity v_z is a fall velocity gained in the drop from MOT to cavity. For short cavities this restriction on v_x is much tighter than the thermal distribution from sub-Doppler cooling to a few μK. Thus the geometrical constraint chops out a flat piece of the thermal distribution, and Φ scales as l_{phys} (or equivalently mode order n) due to this effect.

- 3) The transverse extent of the target mode depends on the cavity waist. The vertical waist affects transit durations but is of little importance for atomic flux. This gives us a factor of w_0 in Φ.

These scalings can tell us how a proposed cavity compares with those already in action. Since the experiment of [51, 52] was the first in our group to incorporate a double-MOT design, we take it as a standard of achievable MOT-to-cavity transfer (with no cooling or catching in the cavity). From several sample data files, I counted a average number of 40 detectable transits per drop in that setup (referred to hereafter as "Lab 11" for the room in which it is located). The Lab 11 cavity has $n = 103$, $w_0 = 23.8\ \mu$m, and $gap = 42.6\ \mu$m. In sum, therefore, we have

$$\Phi \simeq 40(\frac{w_0}{23.8\ \mu m})(\frac{n}{103})(\frac{gap}{42.6\ \mu m}) \tag{7.7}$$

For example, I have calculated the ratio Φ_{Lab11}/Φ for several cavity designs we considered for the active-feedback experiment. Values in Table 7.1 are for mirrors of diameter $d = 1$ mm and symmetric cavities (two identical mirrors). The current active-feedback cavity has length $l_{phys} = 20(\lambda/2)$ and uses $R = 10$ cm mirrors; more atoms in the initial cloud are therefore necessary to ensure ≈ 1 transit per drop.

	R=10 cm mirrors			R=5 cm mirrors		
n (# $\lambda/2$)	$w_0(\mu m)$	gap (μm)	Φ_{Lab11}/Φ	$w_0(\mu m)$	gap (μm)	Φ_{Lab11}/Φ
10	11.2	1.76	530	9.4	none	n/a
11	11.5	2.19	377	9.6	none	n/a
12	11.7	2.61	285	9.8	0.11	8073
13	11.9	3.04	222	10	0.54	1488
14	12.2	3.46	177	10.2	0.96	762
15	12.4	3.89	144	10.4	1.39	482
16	12.6	4.32	120	10.6	1.82	338
17	12.8	4.74	101	10.7	2.24	256
18	13.0	5.17	86	10.9	2.67	199
19	13.1	5.59	75	11.0	3.09	162
20	13.3	6.02	65	11.2	3.52	132
21	13.5	6.45	57	11.3	3.95	111
22	13.6	6.87	51	11.4	4.37	95
23	13.8	7.3	45	11.6	4.8	82
24	13.9	7.72	41	11.7	5.22	71
25	14.1	8.15	36	11.8	5.65	63

Table 7.1: Geometrical properties of cavities as length and mirror curvature are varied.

We keep in mind also that these gaps and fluxes are calculated based on perfect alignment of the cavity mirrors. In reality most cavities may be constructed with a slight relative tilt of the two mirrors, causing the gap to be even smaller at some places around the mirror circumference.

7.3.3 Light in the Cavity

Many aspects of light in a cavity were calculated in detail by Christina Hood and can be read about in her thesis [19]. As mentioned above, when the cavity supports a mode

at a given wavelength, some amount of light actually "lives" in the first few coating layers and this causes the cavity free spectral range to differ from naive expectations. The exact FSR is relevant when designing a cavity to support longitudinal modes at more than one specific wavelength. See the discussion of the active-feedback cavity length in Section 6.3.2 for some treatment of this issue.

For the active-feedback experiment we were constructing a cavity of length $\approx$ 10 μm, and so chose $l_{eff} = 9.2\,\mu$m to get a magic-wavelength FORT two FSR's away. The two-FSR spacing is an attractive scenario because the FORT and probe standing waves overlap perfectly at cavity center, with overlap decreasing as we move away from the center and the standing waves get out of phase. The FORT and probe are completely out of phase halfway from the center and come back into phase at the mirrors (to achieve simultaneous resonance). Alternatively one could imagine the shortest possible cavity to support both cavity QED and the magic-wavelength FORT, with one FSR separating the two wavelengths. In this case the FORT and probe are non-overlapped at cavity center and come into register only at the ends of the cavity. For any future consideration of a cavity along these lines, I note here that the best length is $l_{eff} = 4.5\ \mu m$, which allows a FORT at 936.6 nm one FSR away from 852.359 nm. For different mirror coatings these numbers will change, requiring a new coating model for accurate prediction.

For some purposes the transverse mode spacing of the cavity may also be relevant. Using light in a transverse mode is one way to break cylindrical symmetry of the atom-cavity coupling for purposes of tracking and/or cooling. Alternatively, measurement of transverse mode spacing is one proposed way of determining length for a very short cavity. Geometry (ignoring cavity birefringence but including the contribution of mirror dispersion effects) gives transverse mode spacing

$$\Delta\nu_t = (\frac{c}{2\pi l_{eff}}) * cos^{-1}(1 - l_{eff}/R) \approx \frac{c}{\sqrt{2}\pi l_{eff}^{1/2} R^{1/2}} \tag{7.8}$$

for a symmetric cavity. Measurements I have done on test cavities indicate that this expression reflects reality very closely; cavity length measurement based on this

technique is discussed in the final section of this chapter.

One small but crucial matter is the coupling of light into and out of these cavity modes. A beam with $\simeq 10\,\mu m$ waist inside the cavity expands appreciably on its way through the mirror substrate. A minimal requirement for useful physics is that the beam must cleanly enter/exit the back surface of the substrate! In the "standard" cone mirror design, the back surface is 3 mm tall and 4 mm away from the cavity itself. Beam waist as a function of distance from focus is

$$w(r) = w_0\sqrt{1 + (\frac{r\lambda}{\pi w_0^2})^2} \tag{7.9}$$

and w_0 is set by cavity length and mirror curvature as given above. Extra clearance should be given in case the cavity mode axis is not precisely aligned along the geometric axes of the substrate.

So far this has been a description of modes supported by the cavity, but I now digress to mention the possibility of focusing beams through the cavity from the side. This option can be useful for cooling beams (as in [51, 52]) or for the "classical" pulse in various logic proposals. However, it can become a difficult and even impossible feat as cavity length decreases towards the optical wavelength. Imagine a beam entering through the gap between 1mm diameter cavity mirrors, focused at the cavity axis, and exiting again at the other side. Referring to Equation 7.9, we see that for light of $\lambda = 852nm$, the entering/exiting spot waist $w(r = 0.5mm)$ has a minimum value of $16.5\,\mu m$ achieved with focused waist $w(0) = 12.0\,\mu m$. If we require that a "cleanly" passed beam must have $w(0.5\,mm) \leq gap$, this gives a minimum cavity length of about 33 μm if a side beam is to be used. At shorter cavity lengths light can be focused in and/or allowed to diffract through the cavity volume, but clean beams in *and* out are not a possibility. Smaller mirror diameter would relax this limit, but going below the current 1-mm diameter would require new efforts in mirror cleaning to obtain pristine central regions for the cavity mode with the rough and dirty edge region brought closer in.

7.3.4 Usefulness for Cavity QED

We can begin to assess a cavity's usefulness for cavity QED experiments by looking at relevant rates and at dimensionless parameters that tell us what atom-cavity physics looks like in a given cavity. The rate of coherent atom-cavity coupling is set by the parameter g_0, defined as half the single-photon Rabi frequency. The value of g_0 is simply determined by the electric field strength associated with a single photon in the cavity mode, so g_0 is inversely proportional to the square root of cavity mode volume. Expressing everything in mks units, we have (see, e.g., [27]):

$$\frac{g_0}{2\pi} = (\frac{3c\lambda^2\gamma_{\|}}{2\pi^2 w_0^2 l_{eff}})^{1/2}/2\pi \tag{7.10}$$

where $\gamma_{\|} = 2\gamma_{\perp}$ as given in Chapter 3. This coherent exchange of excitation between atom and cavity must generally be compared with the cavity decay rate κ and the atomic decay rate $\gamma_{\perp}$. If an atom is moving within the cavity, these rates must also be compared with the rate for local coupling to change due to atomic motion; this is represented by $1/\tau$ where τ is a timescale for an atom to fall through the cavity mode or for a trapped atom to oscillate in the direction of interest.

Note that g_0 depends on cavity geometry as

$$g_0 \propto w_0^{-1} l_{eff}^{-1/2} \propto R^{-1/4} l_{phys}^{-1/4} l_{eff}^{-1/2} \approx R^{-1/4} l^{-3/4} \tag{7.11}$$

where the last relation uses $l_{eff} \approx l_{phys} \approx l$. The cavity decay, on the other hand, scales as

$$\frac{\kappa}{2\pi} = \frac{1}{2}\frac{c/2l_{eff}}{Finesse} \propto l_{eff}^{-1}. \tag{7.12}$$

The timescale for atomic motion in general also depends on cavity geometry, but the specific dependence varies with the type of motion being considered.

A simplistic measure of cavity QED properties might be the number of coherent cycles per decay time, which would be estimated by the quantity g_0/β where $\beta = max[\kappa, \gamma, 1/\tau]$. This quantity being much larger than unity is roughly speaking the criterion for strong coupling, the regime in which the atom-cavity physics has

signatures of quantum rather than semiclassical fields at work. However, within the realm of strong coupling there is room for a great deal of variation in what exact phenomena are present and which ones can be observed with good signal-to-noise.

Slightly more specific are the parameters known as the critical atom and photon numbers (N_0, m_0). These parameters are discussed at length in Chapter 2. They represent the number of atoms necessary to significantly alter the cavity response and the number of photons necessary to significantly alter the atomic state. They scale as

$$N_0 = (2\kappa\gamma)/g_0^2 \propto R^{1/2} l^{1/2}, \;\; m_0 = \gamma^2/2g_0^2 \propto R^{1/2} l^{3/2}. \tag{7.13}$$

Both critical parameters should be much less than unity for strong coupling.

Another parameter, the rate of optical information dI/dt, tries to address the question of detecting intracavity physics through measurements of cavity light output (cavity transmission). Roughly, the rate of optical information represents the rate at which transmitted light (or "missing" transmitted light) can give us information about the atom-cavity state. This quantity is also discussed in Chapter 2. The optical information rate scales as

$$dI/dt \sim g_0^2/\kappa \propto R^{-1/2} l^{-1/2}, \tag{7.14}$$

with optical information per atom or per motional cycle carrying an additional factor of the timescale τ.

Whatever the physics or detection strategy, detection efficiency and signal-to-noise ratio will be determined by the ratio of mirror losses (absorption and scatter) to mirror transmission. Most parameters get more impressive as κ gets smaller, so since $\kappa \sim (losses + transmission)$ it is tempting to decrease transmission with ever-more-reflective mirror coatings. However, a given mirror fabrication process will have some minimum achievable loss per mirror (about 1ppm for current REO standards). Decreasing transmission to or below this mark will only result in slight κ decreases, and will greatly erode detection efficiencies as most of the escaping light will be scattered or absorbed rather than transmitted.

Another important consideration along these lines is the use of a two-sided (symmetric) or one-sided (one very good mirror, one more transmissive) cavity. A one-sided cavity offers the prospect of a truly unique, coherent output mode for the cavity. We envision a cavity with single-mirror transmissions (T_1, T_2) and losses (A_1, A_2). Cavities of the current generation have $(T_1 = T_2) > (A_1 = A_2)$, allowing cavity decay to send light out both cavity mirrors equally. A better situation for quantum communication and nonclassical light generation would instead be $T_1 \gg (T_2, A_1, A_2)$. To achieve this situation while remaining in the strong coupling regime, mirrors with very low transmission and losses are necessary, since one typically desires even T_1 about 10 ppm or smaller. Alternatively, output from both sides of a two-sided cavity could be collected and, with some technical difficulty, coherently combined.

7.4 Signal-to-Noise for Atom Orbits

Clearly different figures of merit will apply to each experimental scheme, making different cavity designs suitable for different purposes. As one example, I present some calculations on the signal-to-noise ratio for observing trapped-atom orbits (as in the atom-cavity microscope) in different cavity geometries. As we will see, both intracavity dynamics and detection sensitivity are important factors, and the overall result is not necessarily an intuitively clear function of cavity parameters.

The goal of this set of calculations is to use one experimental paradigm as an example of position sensing considerations in different cavities. The cavities considered are actual cavities used in the group or, in the last two cases, slight variations on existing cavity designs. Cavity geometry and mirror properties set g_0, κ, and w_0. To choose detunings, in each cavity we envision experimental conditions analogous to those in the single-photon triggered-trapping experiments we have already performed. Thus we measure $|\langle a \rangle|^2$ for a probe beam detuned near the lower dressed state of a system where the cavity resonance is red of the atom resonance. The quantities we hold constant (using their values from the atom-cavity microscope) are (1) the cavity detuning from the atom as compared with the cavity width and (2) the probe

detuning *from the lower dressed state resonance* as compared with the cavity width, i.e.,

$$\Delta_{ca}/\kappa = -47/14.2 = -3.3 \tag{7.15}$$

$$\frac{\omega_p - \omega_{|-\rangle}}{\kappa} = \frac{\Delta_{pa} - (\frac{\Delta_{ca}}{2} - \sqrt{g_0^2 + \frac{\Delta_{ca}^2}{4}})}{\kappa} = \frac{-125 - (\frac{-47}{2} - \sqrt{110^2 + \frac{(-47)^2}{4}})}{14.2} = 0.77. \tag{7.16}$$

Then for each cavity I adjust the drive strength to maximize either: (1) the rate of "extra" photons to indicate the atom is present ($\kappa(|\langle a\rangle|^2_{full} - |\langle a\rangle|^2_{empty})$) or (2) the detection bandwidth at which S/N=1 subject to shot noise,

$$B_{S/N=1} = \kappa \frac{(|\langle a\rangle|^2_{full} - |\langle a\rangle|^2_{empty})^2}{|\langle a\rangle|^2_{full}}, \tag{7.17}$$

or (3) the S/N over a radial oscillation period

$$(S/N)_{\tau_{osc}} = \sqrt{\kappa\tau_{osc}} \frac{|\langle a\rangle|^2_{full} - |\langle a\rangle|^2_{empty}}{\sqrt{|\langle a\rangle|^2_{full}}}. \tag{7.18}$$

The results are shown in Table 7.2. In the table I also show the quantity g_0^2/κ for comparison; according to hand-waving arguments, the rate of extra photons should scale roughly with g_0^2/κ, as should the bandwidth for S/N=1. We can see only a rough correspondence, perhaps owing to the *ad hoc* nature of the detuning choices. Certainly as κ becomes roughly equal to the atomic decay $\gamma_\perp/2\pi = 2.6$ MHz, our detuning choices related only to κ become somewhat inappropriate. More interesting is the ACM-like cavity in which κ stayed the same but g_0 was increased by about 20%; sensing ability here clearly does not scale with g_0^2/κ. Evidently for the comparisons I have made here, not only g_0^2/κ but g_0 itself (i.e., in relation to the constant $\gamma_\perp$) is significant.

[In the calculations above I have assumed unit detection efficiency, since we are concerned here with signal-to-noise considerations driven by cavity properties and not by detection strategy. Furthermore, I have considered only the decay out one cavity mirror (by using rate κ rather than 2κ). Thus the relevant detection efficiency

cavity description	cavity $(g_0/2\pi, \kappa/2\pi)$, w_0	$(\Delta_{pa}/2\pi, \Delta_{ca}/2\pi)$
ACM	(110,14.2) MHz, 14.06 μm	(-125,-47) MHz
active-feedback	(130,17) MHz, 13.3 μm	(-148,-56) MHz
Lab 11 (FORT)	(32,4) MHz, 23.8 μm	(-37,-14) MHz
ACM but $l_{eff} = 5.8\,\mu$m	(179,26.7) MHz, 11.7 μm	(-208,-88) MHz
ACM but $R = 5\,cm$ mirrors	(131,14.2) MHz, 11.8 μm	(-146,-47) MHz

g_0^2/κ	$\kappa(\lvert\langle a\rangle\rvert^2_{full} - \lvert\langle a\rangle\rvert^2_{empty})$	$B_{S/N=1}$	$S/N_{\tau_{osc}}$
$5.3 \cdot 10^9$/s	$2.7 \cdot 10^8$/s	15.36 MHz	83
$6.2 \cdot 10^9$/s	$3.3 \cdot 10^8$/s	18.7 MHz	86
$1.6 \cdot 10^9$/s	$0.60 \cdot 10^8$/s	3.10 MHz	70
$7.5 \cdot 10^9$/s	$5.7 \cdot 10^8$/s	32.3 MHz	97
$7.4 \cdot 10^9$/s	$2.7 \cdot 10^8$/s	14.5 MHz	71

Table 7.2: Sensitivity for atomic motion in different cavities (detunings chosen for ACM-like trapping).

is the quantity η' of the discussion in Section 3.7. To include the effect of detection efficiency, $B_{S/N=1}$ must be multiplied by η' and $(S/N)_{\tau_{osc}}$ must be multiplied by $\sqrt{\eta'}$.]

One reason for the lack of g_0^2/κ scaling in Table 7.2 is an inappropriate choice of quantities held constant. The condition of constant Δ_{ca} in Equation 7.15, in particular, helps in maintaining a deep trapping potential but has no particular relevance to sensing. For sensing, it is more reasonable to place the probe a fixed number of linewidths away from the empty cavity, so

$$\frac{\omega_c - \omega_p}{\kappa} = 78/14.2 = 5.5 \tag{7.19}$$

becomes the first condition to apply. To determine ω_a, we can then require that the system's lower dressed state coincide with the probe frequency, giving

$$\omega_{|-\rangle} - \omega_p = (\frac{\omega_a + \omega_c}{2} - \sqrt{g_0^2 + \frac{\Delta_{ca}^2}{4}}) - \omega_p = 0. \tag{7.20}$$

Using these criteria we obtain detunings and sensitivities as shown in Table 7.3 for the same example cavities.

We now see a sensitivity for atom detection that tracks g_0^2/κ much more closely,

cavity description	cavity $(g_0/2\pi, \kappa/2\pi)$, w_0	$(\Delta_{pa}/2\pi, \Delta_{ca}/2\pi)$
ACM	(110,14.2) MHz, 14.06 μm	(-151,-71) MHz
active-feedback	(130,17) MHz, 13.3 μm	(-181,-87.6) MHz
Lab 11 (FORT)	(32,4) MHz, 23.8 μm	(-45.5,-23) MHz
ACM but $l_{eff} = 5.8\,\mu$m	(179,26.7) MHz, 11.7 μm	(-218,-71) MHz
ACM but $R = 5\,cm$ mirrors	(131,14.2) MHz, 11.8 μm	(-221,-143) MHz

g_0^2/κ	$\kappa(\lvert\langle a\rangle\rvert^2_{full} - \lvert\langle a\rangle\rvert^2_{empty})$	$B_{S/N=1}$	$S/N_{\tau_{osc}}$
$5.3 \cdot 10^9$/s	$2.65 \cdot 10^8$/s	14.5 MHz	82
$6.2 \cdot 10^9$/s	$3.28 \cdot 10^8$/s	17.9 MHz	85
$1.6 \cdot 10^9$/s	$0.63 \cdot 10^8$/s	3.2 MHz	73
$7.5 \cdot 10^9$/s	$4.67 \cdot 10^8$/s	23.7 MHz	85
$7.4 \cdot 10^9$/s	$3.34 \cdot 10^8$/s	19.4 MHz	83

Table 7.3: Sensitivity for atomic motion in different cavities (detunings compatible with trapping but chosen for sensing).

though certainly the simple scaling is not reflected perfectly (Figure 7.6). Signal-to-noise over a motional timescale, however, hardly varies at all for the cavities considered here; as the sensitivity for atomic motion improves, the motional timescale for atom-cavity orbits becomes shorter as well.

In general these extra photon rates and S/N ratios are optimized for much higher driving than the fractional photon in the cavity that we use in experiments. Instead, the largest sensitivities occur with one to a few photons in the cavity. A simple estimate shows we should expect this. As in the on-resonance case described in Chapter 2, we argue that the best effect should be seen when the field approximately saturates the atomic response. However, the field is now detuned from the atomic resonance by $\Delta_{pa} \gg \gamma_\perp$. Thus to saturate the atom the cavity must hold not $m_0 = \gamma_\perp^2/2g_0^2$ photons but $m_0(\Delta_{pa}^2/\gamma_\perp^2) = \Delta_{pa}^2/2g_0^2 \approx 1$ photons.

Why do we not go to ten times more driving in the experiment? The answer lies in the fact that Tables 7.2 and 7.3 contain no information about momentum diffusion rates for the atom. We have seen already in Chapter 4 that the semiclassical momentum diffusion is much larger than the fully-quantum calculation for the atom-cavity microscope. As we go to stronger drive strength the dynamics become more semiclassical with respect to the cavity field, and correspondingly the momentum

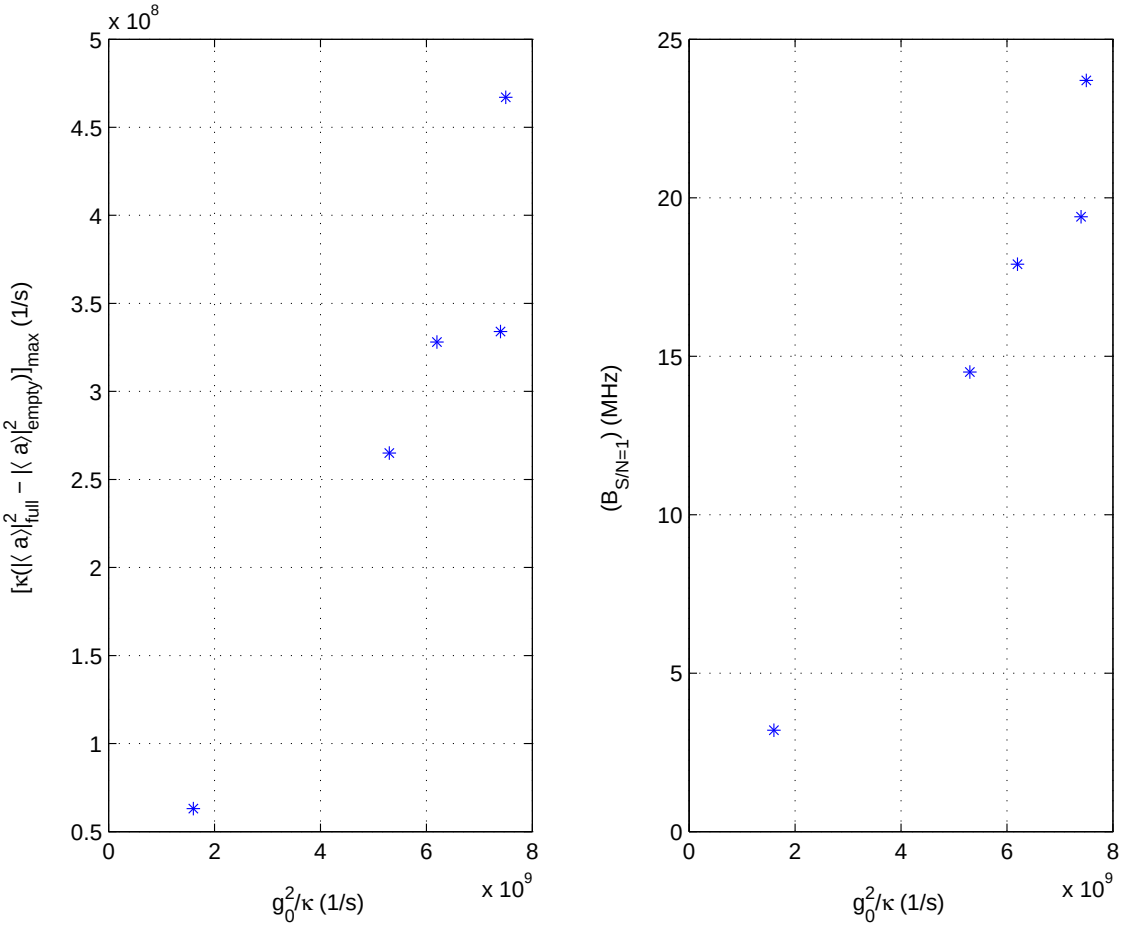

Figure 7.6: Sensitivity for an atom in the cavity as a function of optical information rate (detunings of Table 7.3).

diffusion over oscillation timescales – and over sensing timescales – becomes larger. Thus one has a better bandwidth for sensing atomic position, but the motion itself heats up much faster and there is no quasi-conservative motion to observe.

In an experiment where trapping is provided by some means other than the cavity field, motional timescales are decoupled from the choice of probe settings, in which case the relevant sensitivities are those which do not explicitly include τ_{osc}.

7.4.1 Connection to Position Sensitivity in the ACM

Before leaving the subject of signal-to-noise for detecting trapped atoms, I briefly connect the results of the preceding section with the discussion of position sensitivity

in Chapter 4. In Section 4.3 a hand-waving argument predicted a position sensitivity of 1.0 $nm/\sqrt{Hz}$, notably at odds with the observed sensitivity of 20 $nm/\sqrt{Hz}$. We are now in a position to revisit this estimate and resolve the discrepancy.

The predicted sensitivity arose from an estimate which essentially quoted a rate g_0^2/κ for information on an atom's presence in (or absence from) the cavity mode. This gave sensitivity $S_\rho = \sqrt{\frac{1}{2\eta}}\frac{w_0}{\sqrt{g_0^2/\kappa}}\sqrt{B}$. However, we now have an actual result, from solution of the master equation, for signal-to-noise via our detection method. Thus we may instead observe that, for the atom-cavity microscope, we cannot do better than $S/N = 1$ for resolving an atom's presence in bandwidth $B = 15.36$ MHz. Furthermore, the quantity $dg/d\rho$ estimated as g_0/w_0 in Section 4.3 can be more carefully evaulated, yielding its actual maximum value of $\frac{g_0}{w_0}\sqrt{\frac{2}{e}}$.

Thus we now expect $S_\rho \geq \sqrt{\frac{1}{2\eta'}}\frac{w_0\sqrt{e/2}}{\sqrt{15.36\ \mathrm{MHz}}} = 5.9\ nm/\sqrt{Hz}$. If we allow for a factor of $\sqrt{2}$ degradation in sensitivity due to technical noise, this estimate becomes $8.4\,nm/\sqrt{Hz}$, approaching our measured sensitivity quite reasonably. The remaining discrepancy corresponds to the fact that our probe strength is below the optimal level for sensing, as discussed above.

7.5 Measuring Lengths of Very Short Cavities

In constructing very short cavities, down to the smallest possible Fabry-Perot of length $\lambda/2$, new methods of cavity length measurement will be necessary. We currently measure cavity lengths by tuning a Ti:Sapph laser through a full free spectral range of the cavity, thus finding λ_1 and λ_2 for which $l_{eff} = n(\lambda_1/2) = (n+1)(\lambda_2/2)$. Note n will not be an integer due to mirror dispersion as discussed above. For our current physics cavity of length $l_{eff} = 9.2\,\mu$m, the free spectral range is already 40 nm, and one can see that for very short cavities the free spectral range no longer fits inside the $\sim$150 nm high-transmission region of the mirror coating curve. Alternative length measurements must be employed.

One possibility is to use the transverse mode spacing $\Delta\nu_t$ as a measure of length

for a short cavity. From Equation 7.8 above, we obtain

$$l_{eff} = \frac{c^2}{2\pi^2 R(\Delta\nu_t)^2}. \tag{7.21}$$

A transverse mode spacing may be measured by scanning the cavity length while tuning the laser between *closely spaced* wavelengths λ_1 and λ_2, so the transverse mode spacing on modes at λ_1 can be calibrated against the spacing between the fundamental modes for λ_1 and λ_2. A test cavity with mirror curvature $R = 10cm$ was measured via the usual technique to have $l_{eff} = 11.16\,\mu$m, implying $l_{phys} = 25\lambda/2 = 10.48\,\mu$m since the measurement wavelength was 838nm. For the same cavity a length measurement using the transverse mode technique gave $l_{eff} = 11.06\,\mu$m, agreeing to better than a percent as long as $\Delta\nu_t$ measures l_{eff}. More careful treatment of transverse mode spacing as affected by mirror coating dispersion is necessary if this measurement is to be adopted for very short cavity lengths.

A second candidate measurement technique does not rest on detailed mirror coating calculations but does require the ability to scan the cavity length over a free spectral range with excellent linearity. The simple idea is to calibrate the cavity length scan by observing the scan over a complete free spectral range of some single wavelength, i.e., over $\Delta l = \lambda_1/2$. Then by comparison it is possible to obtain the cavity length difference required between the *same* longitudinal modes for λ_1 and a nearby λ_2. Here $l_{eff,1} = n(\lambda_1/2)$, $l_{eff,2} = n(\lambda_2/2)$, and we measure $l_{eff,1} - l_{eff,2}$ to allow determination of n and thus the overall cavity length. The scan linearity for our current cavity mounts and piezos does not permit a precise measurement via this technique, but it is in principle a promising alternative for the future.